

## METHODS AND APPARATUS FOR PROVIDING DIRECT MEMORY ACCESS CONTROL

### Related Applications

The present application claims the benefit of U.S. Provisional Application Serial No. 60/113,637 entitled "Methods and Apparatus for Providing Direct Memory Access (DMA) Engine" and filed December 23, 1998.

### Field of the Invention

The present invention relates generally to improvements in array processing, and more particularly to advantageous techniques for providing improved mechanisms of data distribution to, and collection from multiple memories often associated with and local to processing elements within an array processor.

### Background of the Invention

Various prior art techniques exist for the transfer of data between system memories or between system memories and I/O devices. Fig. 1 shows a conventional data processing system 100 comprising a host uniprocessor 110, processor local memory 120, direct memory access (DMA) controller 160, system memory 150 which is usually a larger memory store than the processor local memory, having longer access latency, and input/output (I/O) devices 130 and 140.

The DMA controller 160 provides a mechanism for transferring data between processor local memory and system memory or I/O devices concurrent with uniprocessor execution. DMA controllers are sometimes referred to as I/O processors or transfer processors in the literature. System performance is improved since the host uniprocessor can perform computations while the

DMA controller is transferring new input data to the processor local memory and transferring result data to output devices or the system memory. A data transfer is typically specified with the following minimum set of parameters: source address, destination address, and number of data elements to transfer. Addresses are interpreted by the system hardware and uniquely specify I/O devices or memory locations from which data must be read or to which data must be written. Sometimes additional parameters are provided such as element size. One of the limitations of conventional DMA controllers is that address generation capabilities for the data source and data destination are often constrained to be the same. For example, when only a source address, destination address and a transfer count are specified, the implied data access pattern is block-oriented, that is, a sequence of data words from contiguous addresses starting with the source address is copied to a sequence of contiguous addresses starting at the destination address. Array processing presents challenges for data collection and distribution both in terms of addressing flexibility, control and performance. The patterns in which data elements are distributed and collected from processing element local memories can significantly affect the overall performance of the processing system. With the advent of the ManArray architecture it has been recognized that it will be advantageous to have improved techniques for data transfer which provide these capabilities and which are tailored to this new architecture.

### **Summary of the Invention**

As described in detail below, the present invention addresses a variety of advantageous methods and apparatus for improved data transfer control within a data processing system. In particular we provide improved techniques for: distributing data to, and collecting data from an array of processing elements (PEs) in a flexible and efficient manner; and PE address translation which allows data distribution and collection based on PE virtual IDs.

Further aspects of the present invention are related to a virtual-to-physical PE ID translation which works together with a ManArray PE interconnection topology to support a variety of communication models (such as hypercube and mesh) through data placement based upon a PE virtual ID. This result can be accomplished in a DMA controller by translation, through a VID-to-PID lookup table or through combinational logic, where the resulting PID becomes an addressing component on the DMA bus to PE local memories. This result can also be achieved at the PE local memories within the interface logic, where a VID available to the interface logic is compared to a VID presented on the DMA bus. A match at a particular memory interface allows that memory to accept the access. The present invention also addresses the provision of PE addressing modes based on generating data access patterns from logically nested parameterized loops. Varying assignments of loop parameters to nesting level allows flexible data access patterns to be generated. Providing varying mechanisms for updating loop parameters provides greater flexibility for generating complex-periodic access patters, such as select-index modes which provide a table of index-update values which are used when the index loop parameter is updated; select-PE modes which provide a table of bit-vector control values, each of which specifies the PEs to be accessed for an iteration through the "PE update loop" (i.e., the loop which PE update is assigned); and select-index-PE modes which provide both select-index and select-PE update capability and combine to form the most flexible mode for generating complex-periodic data access patterns. Further, the invention addresses the design of a looping mechanism to be reentrant thereby allowing any addressing mode to be restarted after completing a specific number of element transfers, by just loading or reloading a new transfer count and continuing the transfer. This result is accomplished by initializing addressing parameters at instruction load time, and only updating them after a loop exits.

These and other advantages of the present invention will be apparent from the drawings and the Detailed Description which follow.

#### **Brief Description of Drawings**

Fig. 1 shows a conventional data processing system with a DMA controller to support data transfers concurrent with host processor computation;

Fig. 2 illustrates a ManArray DSP with a DMA controller in a representative system in accordance with the present invention;

Fig. 3 illustrates a DMA controller implemented as a multiprocessor, with two transfer controllers, bus connections to a system memory, PE memories and a control bus;

Fig. 4 shows a single transfer controller comprising 4 primary execution units, bus connections and FIFO buffers;

Fig. 5 shows an exemplary format of a transfer type instruction in accordance with the present invention;

Fig. 6 shows an exemplary virtual PE identification to physical PE identification (VID-to-PID) translation;

Fig. 7 shows an exemplary logical implementation of VID-to-PID translation;

Fig. 8 shows an exemplary PEXLAT instruction (“load VID-to-PID table”);

Fig. 9 illustrates a VID-to-PID translation table register, called the PETABLE register in a presently preferred embodiment;

Fig. 10 illustrates a nested logical loop model showing a “BIP” assignment of address components to loops: base (outer), index (middle) and PE VID (inner);

Fig. 11 shows a nested logical loop model with “BPI” assignment of address components to loops: base (outer), PE (middle) and index (inner);

Fig. 12 is a nested logical loop model showing a “PBI” assignment of address components to loops: PE (outer), Base (middle) and Index (inner);

Fig. 13 illustrates an exemplary format for a PE Blockcyclic instruction in accordance with the present invention;

Fig. 14 shows an exemplary transfer result using PE Blockcyclic address mode with BIP loop assignment;

Fig. 15 shows an exemplary transfer result using PE Blockcyclic address mode with BPI loop assignment;

Fig. 16 shows an exemplary transfer result using PE Blockcyclic address mode with PBI loop assignment;

Fig. 17 illustrates an exemplary format for a PE Select-Index transfer instruction in accordance with the present invention;

Fig. 18 shows an exemplary transfer result using a PE Select-Index address mode with BIP loop assignment;

Fig. 19 illustrates an exemplary format for a PE Select-PE transfer instruction in accordance with the present invention;

Fig. 20 shows an exemplary transfer result using a PE Select-PE address mode with BIP loop assignment;

Fig. 21 illustrates an exemplary format for a PE Select-Index-PE transfer instruction in accordance with the present invention; and

Fig. 22 shows an exemplary transfer result using a PE Select-Index -PE address mode with BIP loop assignment.

## Detailed Description

Further details of a presently preferred ManArray core, architecture, and instructions for use in conjunction with the present invention are found in U.S. Patent Application Serial No. 08/885,310 filed June 30, 1997, U.S. Patent Application Serial No. 08/949,122 filed October 10, 1997, U.S. Patent Application Serial No. 09/169,255 filed October 9, 1998, U.S. Patent Application Serial No. 09/169,256 filed October 9, 1998, U.S. Patent Application Serial No. 09/169,072 filed October 9, 1998, U.S. Patent Application Serial No. 09/187,539 filed November 6, 1998, U.S. Patent Application Serial No. 09/205,558 filed December 4, 1998, U.S. Patent Application Serial No. 09/215,081 filed December 18, 1998, U.S. Patent Application Serial No. 09/228,374 filed January 12, 1999 and entitled "Methods and Apparatus to Dynamically Reconfigure the Instruction Pipeline of an Indirect Very Long Instruction Word Scalable Processor", U.S. Patent Application Serial No. 09/238,446 filed January 28, 1999, U.S. Patent Application Serial No. 09/267,570 filed March 12, 1999, U.S. Patent Application Serial No. 09/337,839 filed June 22, 1999, U.S. Patent Application Serial No. 09/350,191 filed July 9, 1999, U.S. Patent Application Serial No. 09/422,015 filed October 21, 1999 entitled "Methods and Apparatus for Abbreviated Instruction and Configurable Processor Architecture", U.S. Patent Application Serial No. 09/432,705 filed November 2, 1999 entitled "Methods and Apparatus for Improved Motion Estimation for Video Encoding", U.S. Patent Application Serial No. \_\_\_\_\_ filed December 23, 1999 entitled "Methods and Apparatus for Providing Data Transfer Control", as well as, Provisional Application Serial No. 60/113,637 entitled "Methods and Apparatus for Providing Direct Memory Access (DMA) Engine" filed December 23, 1998, Provisional Application Serial No. 60/113,555 entitled "Methods and Apparatus Providing Transfer Control" filed December 23, 1998, Provisional Application Serial No.

60/139,946 entitled "Methods and Apparatus for Data Dependent Address Operations and Efficient Variable Length Code Decoding in a VLIW Processor" filed June 18, 1999, Provisional Application Serial No. 60/140,245 entitled "Methods and Apparatus for Generalized Event Detection and Action Specification in a Processor" filed June 21, 1999, Provisional Application Serial No. 60/140,163 entitled "Methods and Apparatus for Improved Efficiency in Pipeline Simulation and Emulation" filed June 21, 1999, Provisional Application Serial No. 60/140,162 entitled "Methods and Apparatus for Initiating and Re-Synchronizing Multi-Cycle SIMD Instructions" filed June 21, 1999, Provisional Application Serial No. 60/140,244 entitled "Methods and Apparatus for Providing One-By-One Manifold Array (1x1 ManArray) Program Context Control" filed June 21, 1999, Provisional Application Serial No. 60/140,325 entitled "Methods and Apparatus for Establishing Port Priority Function in a VLIW Processor" filed June 21, 1999, Provisional Application Serial No. 60/140,425 entitled "Methods and Apparatus for Parallel Processing Utilizing a Manifold Array (ManArray) Architecture and Instruction Syntax" filed June 22, 1999, Provisional Application Serial No. 60/165,337 entitled "Efficient Cosine Transform Implementations on the ManArray Architecture" filed November 12, 1999, and Provisional Application Serial No. \_\_\_\_\_ entitled "Methods and Apparatus for DMA Loading of Very Long Instruction Word Memory" filed December 23, 1999, respectively, all of which are assigned to the assignee of the present invention and incorporated by reference herein in their entirety.

The following definitions of terms are provided as background for the discussion of the invention which follows:

A "transfer" refers to the movement of one or more units of data from a source device (either I/O or memory) to a destination device (I/O or memory).

A data "source" or "destination" refers to a device from which data may be read or to which data may be written which is characterized by a contiguous sequence of one or more addresses, each of which is associated with a data storage element of some unit size. For some data sources and destinations there is a many-to-one mapping of addresses to data element storage locations. For example, an I/O device may be accessed using one of many addresses in a range of addresses, yet it will perform the same operation, such as returning the next data element of a FIFO, for any of them.

A "data access pattern" is a sequence of data source or destination addresses whose relationship to each other is periodic. For example, the sequence of addresses 0, 1, 2, 4, 5, 6, 8, 9, 10,...etc. is a data access pattern. If we look at the differences between successive addresses, we find: 1,1,2, 1,1,2, 1,1,2, ...etc. Every three elements the pattern repeats.

An "address mode" or "addressing mode" refers to a rule that describes a sequence of addresses, usually in terms of one or more parameters. For example, a "block" address mode is described by the rule:  $\text{address}[i] = \text{base\_address} + i$  where  $i = 0, 1, 2, \dots$  etc. and where  $\text{base\_address}$  is a parameter and refers to the starting address of the sequence.

Another example is a "stride" address mode which may be described by the rule:

$$\text{address}[i] = \text{base\_address} + (i \bmod (\text{stride} - \text{hold})) + (i / \text{hold}) * \text{stride}$$

for  $i = 0, 1, 2, \dots$  etc., and where  $\text{base\_address}$ ,  $\text{stride}$  and  $\text{hold}$  are parameters, and where division is integer division in which any remainder is discarded.

An "address generation unit (AGU)" is a hardware module that generates a sequence of addresses (a data access pattern) according to a programmed address mode.

"EOT" means "end-of-transfer" and refers to the state when a transfer execution unit (described in the following text) has completed its most recent transfer instruction by transferring the number of elements specified by the instruction's transfer count field.

The term "host processor" as used in the following description is any processor or device which can write control commands and read status from the DMA controller and/or which can respond to DMA controller messages and signals. In general, a host processor interacts with a DMA controller to control and synchronize the flow of data between devices and memories in the system in such a way as to avoid overrun and underrun conditions at the sources and destinations of data transfers.

The present invention provides a set of flexible addressing modes for supporting efficient data transfers to and from multiple memories, together with methods and apparatus for allowing data accesses to be directed to PEs according to virtual as opposed to physical IDs. This section describes an exemplary DMA controller and a system environment in which the present inventions may be effectively used. The following sections describe PE memory addressing, virtual-to-physical PE ID translation and its purpose, and a set of PE memory addressing modes or "PE addressing modes" which support numerous parallel algorithms with highly efficient data transfer.

Fig. 2 shows an exemplary system 200 illustrating the context in which a ManArray DMA controller 201, in accordance with the present invention, resides. The DMA controller 201 accesses processor local memories 210, 211, 212, 213, 214 and 215 via a DMA Bus 202, 202<sub>1</sub>, 202<sub>2</sub>, 202<sub>3</sub>, 202<sub>4</sub>, 202<sub>5</sub> and memory interface units 205, 206, 207, 208 and 209 to which it is connected. A ManArray DSP 203 also connects to its local memories 210-215 via memory

interface units 205-209. Further details of a presently preferred DSP 203 are found in the above incorporated by reference applications.

In this representative system, the DMA controller also connects to two system busses, a system control bus (SCB) 235 and a system data bus (SDB) 240. The DMA controller is designed to transfer data between devices on the SDB 240, such as a system memory 250 and the DSP 203 local memories 210-215. The SCB 235 is used by an SCB master such as the DSP 203 or a host control processor (HCP) 245 to program the DMA controller 201 with read and write addresses and registers to initiate control operations and read status. The SCB 235 is also used by the DMA controller 201 to send synchronization messages to other SCB bus slaves such as the DSP control registers 225 and a host I/O block 255. Some registers in these slaves can be polled by the DSP and HCP to receive status from the DMA. Alternatively, DMA writes to some of these slave addresses can be programmed to cause interrupts to the DSP and/or HCP allowing DMA controller messages to be handled by interrupt service routines.

Fig. 3 shows a system 300 which illustrates operation of a DMA Controller 301 which may suitably be a multiprocessor specialized to carry out data transfers utilizing one or more transfer controllers 302 and 303. Each transfer controller can operate as an independent processor or work together with other transfer controllers to carry out data transfers. The DMA busses 305 and 310 provide, in the presently preferred embodiment, independent data paths to local memories 320, 321, 322, 323, 324, 325, one for each transfer controller 302 and 303. In addition, each transfer controller is connected to SDB 350 and to SCB 330. Each transfer controller operates as a bus master and a bus slave on both the SCB and SDB. As a bus slave on the SCB, a transfer controller may be accessed by other SCB bus masters in order to read its internal state or to issue control commands. As a bus master on the SCB, a transfer controller

can send synchronization messages to other SCB bus slaves. As a bus master on the SDB, a transfer controller performs data reads and writes from or to system memory or I/O devices which are bus slaves on the SDB. As a bus slave on the SDB, a transfer controller can cooperate with another SDB bus master in a “slave mode” allowing the bus master to read or write data directly from or to its data FIFOs (as discussed further below). It may be noted that the DMA busses 305 and 310, the SDB 350 and the SCB 330 may be implemented in different ways. For example, they may be implemented with varying bus widths, protocols, or the like consistent with the teachings of the present invention.

Fig. 4 shows a system 400 having single transfer controller 401 comprising a set of execution units including an instruction control unit (ICU) 440, a system transfer unit (STU) 402, a core transfer unit (CTU) 408 and an event control unit (ECU) 460. An inbound data queue (IDQ) 405 is a data FIFO buffer which is written with data from an SDB 470 under control of the STU 402. Data is read from the IDQ 405 under control of the CTU 408 to be sent to core memories 430, or sent to the ICU 440 in the case of instruction fetches. An outbound data queue (ODQ) 406 is a data FIFO which is written with data from DMA busses 425 under control of the CTU 408, to be sent to an SDB 470 device or memory under the control of the STU 402. The CTU 408 may also read DMA instructions from a memory attached to the DMA bus, which are forwarded to the ICU 440 for initial decoding. The ECU 460 receives signal inputs from external devices 465, commands from the SCB 450 and instruction data from the ICU 440. It generates output signals 435, 436 and 437 which may be used to generate interrupts on host control processors within the system, and can act as a bus master on the SCB 450 to send synchronization messages to SCB bus slaves.

Each transfer controller within a ManArray DMA controller is designed to fetch its own stream of DMA instructions. DMA instructions are of five basic types: transfer; branch; load; synchronization; and state control. The branch, load, synchronization, and state control types of instructions are collectively referred to as "control instructions", and distinguished from the transfer instructions which actually perform data transfers. DMA instructions are typically of multi-word length and require a variable number of cycles to execute although several control instructions require only a single word to specify. Although the presently preferred embodiment supports multiple DMA instruction types as described in further detail in U.S. Patent Application Serial No. \_\_\_\_\_ entitled "Methods and Apparatus for Providing Data Transfer Control" filed December 23, 1999 and incorporated by reference in its entirety herein, the present invention focuses on instructions and mechanisms which provide for flexible and efficient data transfers to and from multiple memories.

Referring further to system 400 of Fig. 4, transfer-type instructions are dispatched by the ICU for further decoding and execution by the STU 402 and the CTU 408. Transfer instructions have the property that they are fetched and decoded sequentially, in order to load transfer parameters into the appropriate execution unit, but are executed concurrently. The control means for initiating execution of transfer instructions is a flag bit contained in the instruction itself, and is described below.

A "transfer-system-inbound" (TSI) instruction moves data from the SDB 470 to the IDQ 405 and is executed by the STU. A "transfer-core-inbound" (TCI) instruction moves data from the IDQ 405 to the DMA Bus 425 and is executed by the CTU. A "transfer-core-outbound" (TCO) instruction moves data from the DMA Bus 425 to the ODQ 406 and is executed by the CTU. A "transfer-system-outbound" (TSO) instruction moves data from the ODQ 406 to the

SDB 470 and is executed by the STU. Two transfer instructions are required to move data between an SDB system memory and one or more SP or PE local memories on the DMA bus, and both instructions are executed concurrently: a TSI, TCI pair or a TSO, TCO pair.

The address parameter of STU transfer instructions TSI and TSO refers to addresses on the SDB while the address parameter of CTU transfer instructions refers to addresses on the DMA bus to PE and SP local memories.

Fig. 5 shows an exemplary instruction format 500 for transfer instructions. A base opcode field 501 indicates that the instruction is of transfer type. A C/S field 510 indicates the transfer unit (CTU or STU) and I/O field 520 indicates whether the transfer direction is inbound or outbound. The execute ("X") field 550 is a field which, when set to "1", indicates a "start transfer" event, that is, that the transfer should start immediately after loading the transfer instruction. When the "X" field is "0", then the parameters are loaded into the specified unit but the transfer is not initiated. Instruction fetch/decode continues normally until a "start transfer" event occurs. A data type field 530 indicates the size of each element transferred and an address mode 540 refers to the data access pattern which must be generated by the transfer unit. A transfer count 560 indicates the number of data elements of size "data type" which are to be transferred to or from the target memory/device before EOT occurs for that unit. An address parameter 570 specifies the starting address for the transfer. Other parameters 580 may follow the address word of the instruction, depending on the addressing mode used.

While there are six memories 210, 211, 212, 213, 214, and 215 shown in Fig. 2, the PE address modes access only the set of PE memories 210, 211, 212, and 213 in this exemplary ManArray DSP configuration. The address of a data element within PE local memory space is specified with three variables, a PE ID, a base value and an index value. The base and the index

values are summed to form an offset into a PE memory relative to an address 0, the first address of that PE's memory. The address of a PE data element is therefore given by a pair: PE data address = (PE ID, Base + Index).

The ManArray architecture supports a unique interconnection network between processing elements (PEs) which uses PE virtual IDs (VIDs) to support useful single-cycle communication paths, for example, torus or hypercube paths. In some array organizations, the PE's physical and virtual IDs are equal. The VIDs are used in the architecture to specify the pattern for data distribution and collection. When data is distributed according to the pattern established by VID assignment, then efficient inter-PE communication required by the programmer becomes available. As an example, if a programmer needs to establish a hypercube connectivity for a 16 PE ManArray processor, the data will be distributed according to a VID assignment in such a manner that the physical switch connections allow data to be transferred between PEs as though the switch topology were a hypercube even if the switch connections between physical PEs do not support the full hyper-cube interconnect. The present invention describes two approaches whereby the DMA controller can access PE memories according to their VIDs, effectively mapping PE virtual IDs to PE physical IDs (PIDs). The first uses VID-to-PID translation within the CTU of a transfer controller. This translation can be performed either through table-lookup, or through logic permutations on the VID. The second approach associates a VID with a PE by providing a programmable register within the PE or the PE local memory interface unit (LMIU), Fig. 2 205, 206, 207 and 208 which is used by the LMIU logic to "capture" a data access when its VID matches a VID provided on the DMA Bus for each DMA memory access.

## **VID to PID Translation within the DMA Controller**

With this approach, a PE VID-to-PID table is maintained in the DMA controller so that data may be distributed to the ManArray according to a programmer's view of the array. In the preferred embodiment, this table is maintained in the CTU of each transfer controller. Fig. 6 shows an exemplary mapping table 600 of VID into PID for a four PE system, such as a ManArray 2x2 system. The VIDs are in column 602 on the left and their corresponding PIDs are shown in column 604 on the right. An example of a table lookup implementation of the mapping of Fig. 6 is illustrated logically as system 700 of Fig. 7. In the presently preferred embodiment, a translation table 710 is stored in the CTU of a transfer controller. A CTU transfer instruction 705 (TCI or TCO) specifies a starting address 775 which is used by AGU 770 to generate an initial VID 720. The VID 720 controls the selection of one of the elements of the VID-to-PID lookup table 710 through multiplexer 715 which is then sent to a DMA Bus 740 as the PE ID component of the PE address. The numbers on the multiplexer 715 indicate the VID value which must be applied to select the corresponding input. Successive VIDs are generated by the AGU 770, possibly in a recursive fashion as shown by feedback 708. At the same time, the AGU 770 generates a sequence of PE memory offsets 730, also possibly using recursive feedback 755. The PE memory offset 750 is also sent to the DMA bus as a second component of a PE address. Logic in the local memory interface units (LMIUs) is used to compare the PE ID sent on the DMA bus to a stored PID (hard-coded) for any DMA bus access. If this matches, then the LMIU accepts the access and accepts write data or returns read data.

The approach of Fig. 7 has the advantage that all mappings of PE VIDs to PIDs are supported. With larger numbers of PE local memories, the register or memory space required to store this table grows. For example, a 16 PE memory system requires 64 bits of register or

memory space to store the PIDs. An alternative approach to table lookup-based translation is to provide logic which performs a subset of all VID-to-PID mappings. This translation logic would also be parameterized, but would require significantly fewer bits to configure. As a simple example, let the PID be formed by complementing any bit of the VID. If the PID and VID require 4 bits to represent the needed IDs, say for a 16 PE system, then a four bit “translation vector” (XVEC) must be stored to configure the translation rather than the 64 bits for table lookup. The PID is obtained from the VID by the following:  $\text{PID} = \text{VID} \text{ xor } \text{XVEC}$ . That is, each bit of VID is exclusive-or'd with the corresponding bit of XVEC. The set of PIDs resulting from applying this operation to each VID constitutes the mapping. Obviously, the number of mappings available is far fewer than with a table lookup approach, but for systems with a large number of PE memories, only a few mappings may be required to support the desired communication patterns.

In the presently preferred embodiment, a lookup table is used to perform the VID-to-PID translation. Two approaches are provided for initializing the translation table. The first is through a DMA instruction 800, shown in Fig. 8. When executed, DMA instruction 800 loads a PETABLE register 900 which is illustrated in Fig. 9. The second approach is through a direct write of the PETABLE register 900 via the SCB.

#### **PE Virtual IDs Stored in Local Memory Interface Units**

The second approach to directing data access according to PE VID relies on distributing the PE VIDs to each PE local memory interface unit (LMIU). The VID for each PE might reside in a register either in the PE itself or in its LMIU. In this case, there is no translation table or logic in the DMA lane controllers. In common with the preceding approach, there is a PE ID component of the DMA bus which is driven by the transfer controllers and used by the LMIUs

to compare for a match with the locally visible PE VID. When a match is detected in a PE, then it accepts the access which may be either a write or a read request. Means for updating the VIDs stored locally in the LMIUs may be provided through the use of registers visible in the PE register address space, or through a PE instruction which broadcasts the table to all PEs, who then select their VID using their hard-coded PID stored locally. This approach has advantages when VIDs are used for other purposes than just data distribution and collection by a DMA controller.

### **CTU Addressing Modes**

A CTU 408 shown in Fig. 4 supports a basic set of address modes which may be used to target memories associated with each PE or SP individually. These address modes include single-address, block, stride and circular modes. These addressing modes will not be described in detail herein, but are a common set of addressing modes used for many uniprocessor applications. In addition to these address modes, the CTU 408 provides a set of “PE address modes” which allow data to be distributed across or collected from multiple PE memories in a variety of patterns. These address modes are based on a software model of address generation based on parameterizable loops, which is then implemented in hardware.

### **Flexible PE Addressing Modes through Parameterizable Logical Loops**

Many algorithms which are distributed across multiple PEs require complex data access patterns to achieve peak efficiency. The basis for our loop-based PE addressing modes is a logical view of data access consisting of a set of nested loops in which one component of the PE memory address is assigned to be updated at the end of each loop. As stated above, a PE memory address consists of three components called “address components”, a PE virtual ID (VID), a base value (Base) and an index value (Index). This model requires the following: a

mechanism for assigning address components to logical loops; a mechanism for initializing address components; and a mechanism for updating address components; and a mechanism for indicating a loop's exit condition.

Assignment of an address component to a loop specifies the order in which the three address components are updated. In an embodiment which uses a three-loop model, there are six possible orders for updating address components (i.e. six ways to re-order VID, Base and Index). The base and index components are defined to be ordered in this embodiment so that the index is always updated prior to the base, which reduces the number of possible orderings to three, since base and index are summed to form an offset into PE memory, allowing loop assignments that update the base before the index is redundant. An exemplary loop assignment is: update VID on inner loop; update index on middle loop; and update base on outer loop.

Thus, as PE addresses are generated, the VID component updates first (inner loop). When all VIDs have been used (VID loop exit condition has been reached), then the VID is reinitialized, the index is updated, and the VID loop is reentered. This looping continues until the number of index updates is exhausted (Index loop exit condition has been reached) at which point the index is reinitialized, the base is updated, the index loop is reentered, then the VID loop is reentered. This further looping continues until the transfer count is exhausted.

Updating an address component is performed by selecting a new value for the component either based on the old value (e.g.  $\text{new} = \text{old} + 1$ ) or by some other means, such as by table lookup. A loop exit condition specifies what causes the loop to exit to the next-most outer loop in the model.

In summary, three different aspects of loop control are used to vary the sequence in which PE memories may be accessed. These are:

- (1) Rearranging the order of assignment of address components to logical loops,
- (2) Varying the method for updating the address components, and
- (3) Varying the loop termination conditions.

Figs. 10, 11 and 12 show logical representations or processes 1000, 1100 and 1200, respectively, of preferred assignments of address parameters (PE VID, Base and Index) to logical loops. In the nomenclature used in Figs. 10, 11 and 12, the term “PE” refers to the PE VID component of a PE address. In Fig. 10, the address components are assigned in “Base, Index, PE” (BIP) ordering. This means that the PE is updated in the innermost loop, the index parameter is updated in the “middle” loop and the base parameter is updated in the “outer” loop. In Fig. 11, the loop assignments are in a “Base, PE, Index” (BPI) ordering, and in Fig. 12, the loop assignments are in a “PE, Base, Index” (PBI) ordering.

Fig. 10 shows a logical representation 1000 of the nested loop model in which the PE VID is updated in an inner loop 1030, the index is updated in a middle loop 1020, and the base is updated in an outer loop 1010. A fourth loop 1005 which encompasses the other three loops indicates that the other loops are continued until the number of data elements specified in the transfer instruction have been accessed. Associated with each loop is a condition for loop exit 1010, 1020 or 1030, respectively, where the “!” character represents a logical NOT. Also associated with each loop is a mechanism 1060, 1070 or 1077, respectively, for updating the loop address parameter and for testing the updated value to indicate whether the exit condition for that loop has become TRUE. Prior to starting any loop is an address initialization block 1002 which sets the starting values of each address component (PE, Base and Index). The data transfer implemented by Fig. 10 will cause PEs to be accessed first until an “exit PE loop” condition has become true (PELoopComplete is TRUE), at which point the PE loop exits and the PE parameter

is reinitialized in step 1065. The index parameter is then updated and tested for its terminal condition in step 1070. If the index parameter's terminal condition has not become TRUE, then the PE loop is reentered. When the index parameter's terminal condition becomes TRUE, the index loop is exited, the index parameter is reinitialized in step 1075 and the base parameter is updated and tested for a terminal condition in step 1080. If the base parameter terminal condition has not been reached, then the index and PE loops are reentered and executed until either all data items have been accessed (transfer count specified in the transfer instruction becomes zero) or the index loop is terminated again. When BaseLoopComplete becomes TRUE, the base value is reinitialized in step 1085 and the loops are reentered again.

Figs. 11 and 12 show nested logical loops or processes 1100 and 1200 corresponding to "BPI" access (index is updated first, followed by PE, followed by base) and "PBI" access (Index is updated first, followed by Base, then lastly PE) respectively.

The following aspects of the loop formulation are noted. When the requested number of accesses are made (TC in Figs. 10-12) then all loops are exited immediately, leaving all address and loop control variables in their current states. By using logical "while" loops and reinitializing a loop only at its exit, it is possible to reenter the loops and continue a transfer after "terminal count" (TC) addresses have been accessed. This capability is used in this invention to allow transfers to be restarted so that the addressing continues as though it would if the transfer count had not been exhausted. For further details of such transfers see U.S. Application Serial No. \_\_\_\_\_ filed December 23, 1999 entitled "Methods and Apparatus for Providing Data Transfer Control" which is incorporated by reference in its entirety herein.

The functions used to update an address (see UpdateAddress() in Fig. 10 steps 1060, 1070 and 1077; in Fig. 11 steps 1160, 1170 and 1177; and in Fig. 12 steps 1260, 1270 and 1277)

may update the address using a constant increment value, or a value extracted from a table, or use a selection mechanism based on a bit vector. While other UpdateAddress() functions might be supported, those listed are supported in the presently preferred embodiment.

The function used to update the loop control variable, UpdateLoopControl( ), may be performed as part of the address update or as a separate operation as shown in Figs. 10-12. This operation is used to update variables which control loop termination. In the preferred embodiment, the control variables are counters or special logical functions consisting of priority encoders and counter blocks.

The function used to check for loop termination simply tests the loop termination variable for an end of loop condition. This condition may be a particular count value or the state of a mask register.

The initialization of address parameters (see Initialize() function: Fig. 10 1002, Fig. 11 1102, and Fig. 12 1202) does not necessarily occur each time a transfer is started. In the preferred embodiment, this initialization occurs only when a transfer instruction is decoded and parameters are loaded into CTU registers in the case of PE addressing modes or STU registers.

The following discussion addresses instruction formats and describes PE addressing modes for one embodiment of the invention. It will be recognized other instruction encodings may be used consistent with the teachings of the present invention. In the preferred embodiment, a transfer controller reads transfer instructions from a local memory and decodes them. Transfer instructions come in two types, those for the STU and those for the CTU. The STU transfer instructions specify the addressing mode and transfer count for accesses to the system data bus while CTU transfer instructions specify the addressing mode and transfer count for accesses to the DMA bus and all SP and PE memories. The instruction formats addressed below are only

those instructions which control special PE memory addressing for the CTU. Instruction mnemonics are used to indicate the instruction type and addressing mode. “TCI” stands for “transfer, core- inbound”, while “TCO” stands for “transfer, core-outbound”. “TCx” stands for either TCI or TCO. The following PE addressing modes are described as illustrative of the present invention: PE Block-Cyclic, PE Select-Index, PE Select-PE, and PE Select-Index-PE.

### **PE Block-Cyclic Addressing**

PE blockcyclic addressing provides the basic framework for all of the PE addressing modes. A Loop parameter specifies the assignment of address components to loops: BIP, BPI, or PBI. Fig. 13 shows an exemplary format 1300 which defines the parameters for a PE Blockcyclic transfer instruction executed by the CTU. As an example, if we are given:

An inbound sequence of 16 data elements with values 0,1,2,3,...15;  
PETABLE setting of 0x000000E4 (no translation of PE IDs);  
TSI.block instruction in the STU (reading the 16 values from system memory); and  
TCI.blockcyclic instruction in the CTU with PE count = 4, Base Update = 8, Base Count=2 (used for PBI mode only), Index Update = 2, Index Count = 2, then the resulting data in the PE memories 1400 after the transfer are shown in Fig. 14 for BIP loop assignment. Fig. 15 shows resulting data 1500 for BPI loop assignment. Fig. 16 shows resulting data 1600 for PBI loop assignment.

### **PE Select-Index Addressing**

The operation of the PE select-index address mode is similar to the PE blockcyclic address mode except that rather than updating the index component of the address by adding a constant to it, the instruction specifies a table of index update values which are used sequentially

OPERATION CODE

to update the index. Fig. 17 shows an exemplary instruction format 1700 for the PE select-index instruction.

An index select parameter allows finer-grained control over a sequence of index values to be accessed. In the example, this is done using a table of eight 4-bit index-update (IU) values. Each time the index loop is updated, an IU value is added to the effective address. These update values are accessed from the table sequentially starting from IU0 for IUCount updates. After IUCount updates, the index update loop is complete and the next outer loop (B or P) is activated. On the next entry of the index loop, IU values are accessed starting at the beginning of the table. Fig. 18 shows an exemplary data access table 1800 illustrating data access using the PE select-index instruction.

### PE Select-PE Addressing

The operation of the PE Select-PE address mode is similar to the PE blockcyclic address mode except that rather than updating the PE VID component of the address by adding 1 to it, the instruction specifies a table of bit vectors, where each bit vector specifies the PE's to select for access. A bit set to "1" in a bit vector indicates, by its bit position, the VID of the PE to access. Bits in each bit vector are scanned from right to left (least to most significant when viewed in a first instruction format such as instruction format 1900 of Fig. 19). When there are no more "1" bits in a vector, the PE loop exits. The next iteration of the loop uses the next bit vector in the table. Fig. 19 shows an exemplary instruction format 1900, and Fig. 20 shows an exemplary transfer data access table 2000 for a transfer using this instruction.

The PE select fields together with the use of the PE translate table allow out of order access to PEs across multiple passes through them.

### **PE Select-Index-PE Addressing**

This addressing mode combines both select-index and select-PE addressing . An exemplary instruction format 2100 is shown in Fig. 21. This form of addressing provides for complex-periodic data access patterns. An exemplary access pattern table 2200 for the PE-select-index-PE address mode is shown in Fig. 22.